

Small H -coloring problems for bounded degree digraphs

Pavol Hell *

School of Computing Science,
Simon Fraser University,
Burnaby, B.C., Canada - V5A 1S6,
pavol@cs.sfu.ca

Auroshish Mishra‡

Department of Computer Science,
Cornell University,
Ithaca, NY, USA - 14853,
auroshish@cs.cornell.edu

Abstract

An NP-complete coloring or homomorphism problem may become polynomial time solvable when restricted to graphs with degrees bounded by a small number, but remain NP-complete if the bound is higher. For instance, 3-colorability of graphs with degrees bounded by 3 can be decided by Brooks' theorem, while for graphs with degrees bounded by 4, the 3-colorability problem is NP-complete. We investigate an analogous phenomenon for digraphs, focusing on the three smallest digraphs H with NP-complete H -colorability problems. It turns out that in all three cases the H -coloring problem is polynomial time solvable for digraphs with degree bounds $\Delta^+ \leq 1$, $\Delta^- \leq 2$ (or $\Delta^+ \leq 2$, $\Delta^- \leq 1$). On the other hand with degree bounds $\Delta^+ \leq 2$, $\Delta^- \leq 2$, all three problems are again NP-complete. A conjecture proposed for graphs H by Feder, Hell and Huang states that any variant of the H -coloring problem which is NP-complete without degree constraints is also NP-complete with degree constraints, provided the degree bounds are high enough. Our study is the first confirmation that the conjecture may also apply to digraphs.

1 Introduction

Graph coloring problems arise naturally in several contexts of both theoretical and applied nature. Be it scheduling events, solving pattern matching problems, or allocating registers to processes, many constraint satisfaction problems can be easily modeled in the graph coloring setting. A slightly more general problem is that of H -coloring.

Consider a fixed graph H . A *homomorphism* $f : G \rightarrow H$ is a mapping $f : V(G) \rightarrow V(H)$ such that $f(u)f(v)$ is an edge of H for each edge uv of G . An H -*coloring* of G is a homomorphism $G \rightarrow H$. The definition is formally the same for digraphs G, H : since the edges are directed, the homomorphism f is a mapping that preserves both the edges and their direction.

The H -*coloring problem* is the decision problem which asks the following question:

Given a graph G , is it possible to H -color G ?

The complexity of the H -coloring problem for graphs H has been widely studied [10]. It is shown in [9] that the H -coloring problem is polynomial time solvable if H is bipartite or contains a loop, and is NP-complete otherwise. A number of variants of this basic family of problems have been considered, and in [8] the authors have set up a framework for these variants. One of the

*Partially supported by a NSERC discovery grant

†The facilities of the IRMACS center at SFU are gratefully acknowledged, where most of this research was done

‡Supported by the MITACS Globalink Internship Program for undergraduate students

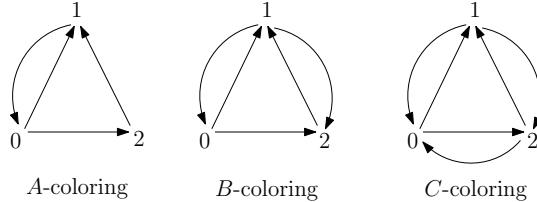


Figure 1: The digraphs A , B , and C

parameters in this framework is a restriction to graphs with a given upper bound on the vertex degrees. For graphs with degrees bounded by 3, some H -coloring problems that are NP-complete in general, become polynomial time solvable. However, these problems tend to be NP-complete again when the degree bound is 4. Such is, for instance, the situation with $H = K_3$, i.e. with 3-colorings. The theorem of Brooks ensures that a connected graph with degrees at most 3, other than K_4 , is 3-colorable [3]. Thus 3-colorability of such graphs is decidable in polynomial time. However, it is known that the 3-colorability of graphs with degrees at most 4 is NP-complete [11]. (Our results also imply this fact, see Section 4.) Based on additional evidence of this kind, Feder, Hell and Huang conjectured that any variant of the H -coloring problem which is NP-complete without degree constraints is also NP-complete with degree constraints, provided the degree bound is high enough [4].

In contrast, for digraphs H , the boundary between easy and hard H -coloring problems is not well-understood. Some partial results have been published in [12, 1, 7, 2] and in several other papers, but it is still not known whether all directed H -coloring problems are polynomial or NP-complete. This statement is in fact equivalent to the well-known Dichotomy Conjecture of CSP [5]. Even less is known about the complexity of H -coloring for digraphs with bounded degrees, and its relation to the conjecture of Feder, Hell, and Huang mentioned above. Of course, K_3 can be viewed as a symmetric digraph, namely the digraph C in Figure 1, and in this context C -coloring of a digraph G is precisely a 3-coloring of the underlying undirected graph of G . The digraphs A , B , and C , in Figure 1, are the smallest three digraphs H with NP-complete H -coloring problems. All other digraphs H with three vertices have polynomial-time solvable H -coloring problems [1, 12]. (All digraphs H with fewer than three vertices have H -coloring problems that can be solved in polynomial time by 2-SAT [10].) We investigate these first three interesting cases. We confirm that the H -coloring problems for these three digraphs, $H = A, B$, and C , are polynomial time solvable for low degree bounds, and become NP-complete again when the degree bounds are high enough.

Interestingly, we can determine very precisely the highest allowable degrees that ensures the polynomiality of these problems. Let Δ^+ denote the maximum out-degree and Δ^- the maximum in-degree in a digraph. We show that digraphs with $\Delta^+ \leq 1$, $\Delta^- \leq 2$, or with $\Delta^+ \leq 2$, $\Delta^- \leq 1$, have polynomial time A -coloring, B -coloring, and C -coloring algorithms, but for digraphs with $\Delta^+ \leq 2$, $\Delta^- \leq 2$, all three problems are again NP-complete. The NP-completeness of C -coloring will imply the theorem of Holyer [11] mentioned earlier.

For the polynomial cases, we provide algorithms for the more general problem of *list H -coloring*.

2 General Digraphs

In this section, we will present new NP-completeness proofs for the A -coloring, B -coloring, and C -coloring of general digraphs, i.e. without any degree restrictions. It is a folklore result that these three problems are NP-complete [1, 12]. The proofs we provide here will facilitate extensions to

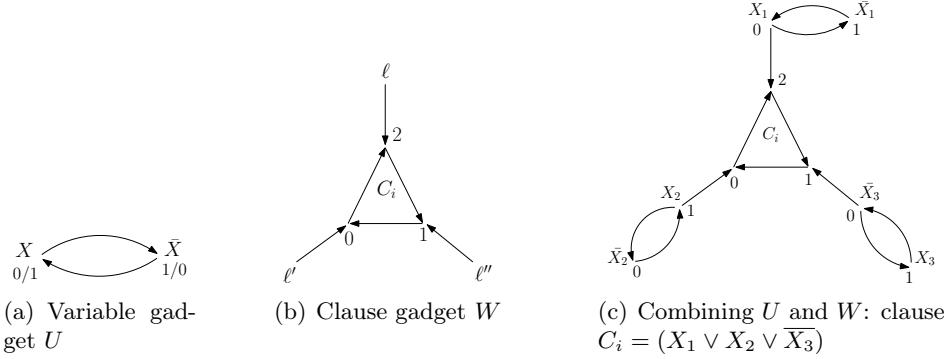


Figure 2: NP-completeness of A -coloring

graphs with bounded degrees.

As is standard in such proofs, we will have a *variable gadget* for each variable and a *clause gadget* for each clause. We will reduce the problem 1-in-3-SAT to the A - and B -coloring problems, i.e., for each formula ϕ produce a graph G_ϕ in such a way that ϕ is satisfiable if and only if G_ϕ has an H -coloring. For the C -coloring problem, we will show a reduction from the 3-SAT problem.

We begin with A -coloring. We construct a reduction from 1-in-3-SAT as follows. Let ϕ be a given 3-CNF formula. Consider the variable gadget U shown in Figure 2(a). For each variable X in ϕ , we have one such gadget, with the two endpoints of U corresponding to the variable X and its negation \bar{X} .

Figure 2(b) depicts the clause gadget W , for a generic clause $C_i = (\ell \vee \ell' \vee \ell'')$ with three literals ℓ, ℓ', ℓ'' . The vertices $\ell, \ell', \ell'', \ell'''$ are identified with the same literals of the corresponding variable gadgets, producing a graph G_ϕ . Note that an endpoint of a variable gadget has an outgoing edge for each occurrence of the corresponding variable (or its negation) in a clause of ϕ .

The graph A has only one digon, with the vertices 0 and 1. Thus, there are exactly two ways of A -coloring the variable gadget U , with X colored either 0 or 1 and \bar{X} colored 1 or 0, respectively, as depicted in Figure 2(a). The inner 3-cycle of the clause gadget W requires three different colors, and in fact the colors must appear in the cyclic order 0, 2, 1, as depicted in the figure, up to the symmetry of W . Consider now the possible colors of the vertices ℓ, ℓ', ℓ'' (up to symmetry):

- (0,0,0) : as 00 is not an edge, this would not allow a 0 on the inner cycle;
- (1,1,1) : as 11 is not an edge, this would not allow a 1 on the inner cycle;
- (0,1,1) : as 10 is the only edge from 1, this would force two 0's on the inner cycle;
- (0,0,1) : Figure 2(c) depicts the unique A -coloring

It follows that in any A -coloring of G_ϕ , exactly one literal in each clause is colored by 1. Thus if we assign the value True to the literals colored 1 and value False to the literals colored 0, we obtain a satisfying truth assignment for ϕ . Conversely, if we start with a truth assignment and color all vertices of the variable gadgets 0 if the literal is False and 1 if the literal is True, there is, in each clause, exactly one literal that is colored 1, and so the above analysis shows that the colors can be extended to the inner cycles of all the clause gadgets. In conclusion, ϕ is satisfiable if and only if G_ϕ is A -colorable.

Since the A -coloring problem is clearly in NP, we have the following theorem.

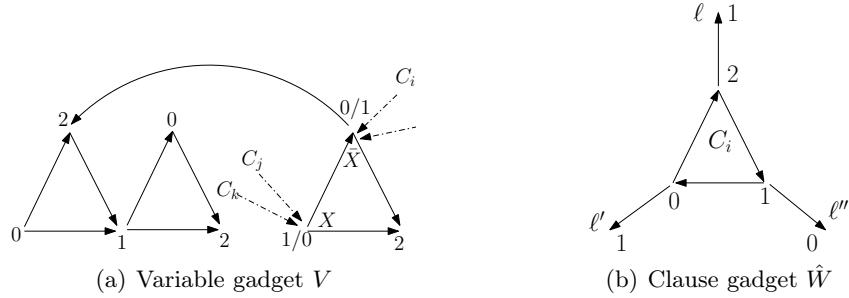


Figure 3: NP-completeness of B -coloring

Theorem 1. *The A -coloring problem is NP-complete.*

A similar analysis applies for B -coloring. The variable gadget V is shown in Figure 3(a). It is more complex than for A -coloring, but there are still two vertices corresponding to the variable X and its negation \bar{X} . It is important to observe that V can be B -colored so that X, \bar{X} are colored by 0, 1 or by 1, 0, but no other pair of colors. To see this, note that the vertex lying in the two adjoining triangles must be colored 1 in any B -coloring of V , thus forcing the colors of the adjoining triangles of V as depicted in the figure. Therefore, the color of \bar{X} is either 0 or 1 (since it has an out-neighbor colored 2). On the other hand, the color of X cannot be 2, since X has two adjacent out-neighbors. Thus coloring \bar{X} by 0 forces X to be colored 1 and vice versa, and no other pair of colors for X, \bar{X} is possible.

Figure 3(b) depicts the clause gadget \hat{W} for $C_i = (\ell \vee \ell' \vee \ell'')$ of the 3-CNF formula. This gadget is similar to the gadget W used for A coloring, except the edges now point away from the center cycle, into the variable gadgets. In any B -coloring, the inner cycle requires three different colors. Again, colors in the order of 0, 2, 1, as depicted, give the only B -coloring for such a cycle (up to symmetry). The vertices ℓ, ℓ', ℓ'' are identified with the corresponding literals. They can only be colored 0 or 1. This again gives rise to four possibilities for the colors of ℓ, ℓ', ℓ'' (up to symmetry), but here only (0,1,1) is a valid choice:

1. (0,0,0) : as 00 is not an edge, this would not allow a 0 on the inner cycle;
2. (1,1,1) : as 11 is not an edge, this would not allow a 1 on the inner cycle;
3. (0,0,1) : as 20 is not an edge, this would force two 1's on the inner cycle;
4. (0,1,1) : Figure 3(b) depicts the unique B -coloring.

Now, if we assign a truth value of True to the literals that are colored 0, and False to the literals that are colored 1, we have a satisfying assignment for the 1-in-3-SAT problem, and vice-versa. Thus, we have the following theorem.

Theorem 2. *The B -coloring problem is NP-complete.*

Finally, we handle the case of C -coloring.

Recall that a C -coloring of a digraph G is precisely a 3-coloring of the underlying undirected graph of G . Thus we can use the proof from [6], which constructs a reduction from 3-SAT using the clause gadget depicted in Figure 4(b). The variable gadget depicted in Figure 4(a), and the identification of the literals ℓ, ℓ', ℓ'' with the same literals of the appropriate clause gadgets of G_ϕ are made as before. Assuming the top vertex of the clause gadget is colored 1 as depicted, the

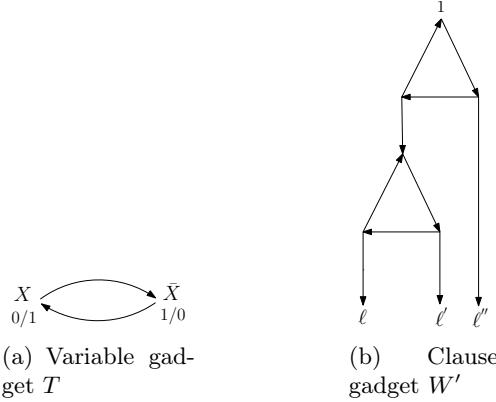


Figure 4: NP-completeness of B -coloring

three vertices labelled ℓ, ℓ', ℓ'' (corresponding to the literals) cannot be colored by $0, 0, 0$, but can be colored by any other combination of $0, 1$. Now it is easy to see that an instance ϕ of 3-SAT is satisfiable if and only if the corresponding digraph G_ϕ (i.e., its underlying undirected graph) is 3-colorable cf. [6].

Theorem 3. *The C -coloring problem is NP-complete.*

3 Bounded Degree Digraphs

We observed in the previous section that the A -, B -, and C -coloring problems are NP-complete for general digraphs. Now, we investigate the same problems restricted to the class of digraphs with bounded in- and out-degrees.

We begin by presenting a polynomial-time H -coloring algorithm for digraphs H with degrees $\Delta^+ \leq 2$, $\Delta^- \leq 1$, or $\Delta^+ \leq 1$, $\Delta^- \leq 2$.

Let H be any fixed digraph. A digraph G with degrees $\Delta^+ \leq 2$, $\Delta^- \leq 1$, or $\Delta^+ \leq 1$, $\Delta^- \leq 2$ is either a binary tree or a directed cycle with binary trees attached at vertices of the cycle (see Figure 5(c) illustrating the case $\Delta^+ \leq 2$, $\Delta^- \leq 1$). Oriented cycles and multiple directed cycles in a single component are not possible since they would violate the degree requirement at some vertex.

As mentioned earlier, we actually solve a more general problem called list H -coloring. To begin with, let us introduce it formally. Given a fixed digraph H , the problem of *list H -coloring* is the following:

Given an input graph G , and for each $v \in V(G)$, a list $L(v) \subseteq V(H)$, is there a homomorphism $f : G \rightarrow H$ such that $f(v) \in L(v)$ for all vertices $v \in V(G)$?

We note that by setting all lists $L(v) = V(H)$, we reduce the H -coloring problem to the list H -coloring problem. Thus, solving the list H -coloring problem in polynomial time also solves the H -coloring problem in polynomial time.

Algorithm. *Arc Consistency* is a basic technique from artificial intelligence, useful for solving list homomorphism problems [10]. In particular, it consists of considering an edge uv of the input graph G and reducing the lists $L(u), L(v)$ so that for each $x \in L(u)$ there exists a $y \in L(v)$ with $xy \in E(H)$, and for each $z \in L(v)$ there exists a $w \in L(u)$ with $wz \in E(H)$. This kind of constraint

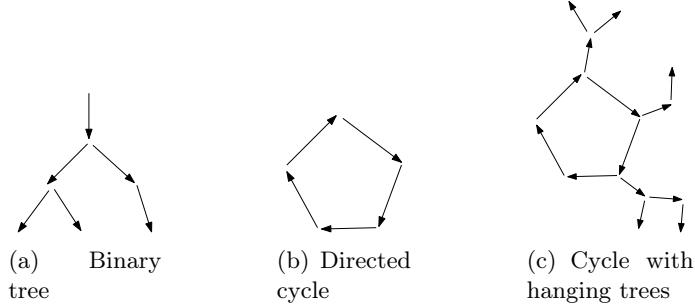


Figure 5: Bounded digraphs: $\Delta^+ \leq 2$, $\Delta^- \leq 1$

propagation will make the entire graph arc consistent by repeating this removal as long as possible, while the lists in G are changing. Since H is fixed, after linearly many updates we obtain final lists satisfying all these constraints. It is easy to observe that if there is at least one empty final list, there is no homomorphism from G to H satisfying the original lists.

When G is a tree and the final lists are all non-empty, there is a list homomorphism $G \rightarrow H$ satisfying the original lists - simply choose one element from the final list of one fixed vertex and propagate this choice to all other vertices using the fact that both constraints were satisfied. This is not true, however, in general. For instance, when H is the 2-cycle with edges 01, 10 and G is a directed cycle of odd length (see Figure 5(b)) with all lists equal to $\{0, 1\}$, then the conditions are satisfied but there is no homomorphism. If we choose one vertex of the 5-cycle and map it to 0, say, then the constraints will propagate around the cycle and arrive requiring the choice 1 at the chosen vertex (and vice versa). However, if at least one vertex of a cycle has a list of size one, the constraint propagation around the cycle will work properly.

With this knowledge in mind, we propose our list H -coloring algorithm:

1. Run an arc-consistency algorithm over all the edges of G .
2. If some vertex has an empty list after the arc-consistency process, then, there is no list homomorphism.
3. Else, find an H -coloring for each weak component of G separately.

If the component is a tree, choose one vertex and map it to one member of its list and propagate this choice to the entire component.

If the component is a directed cycle, choose one vertex u , and consider $|L(u)|$ subproblems, in which the list $L(u)$ is reduced to a single vertex of H , i.e. for each of the possible $L(u)$ choices. Perform constraint propagation through the component based on this choice. If at least one choice leads to a homomorphism, then we have a solution, otherwise there is no solution.

If the component has a directed cycle with hanging trees, follow the procedure to find a homomorphism for the directed cycle first. Then, extend it to find a homomorphism for each of the hanging trees.

Lemma 1. *Let H be a fixed digraph. The list H -coloring problem is polynomial time solvable for digraphs with $\Delta^+ \leq 2$, $\Delta^- \leq 1$, or $\Delta^+ \leq 1$, $\Delta^- \leq 2$.*

In particular, this applies to $H = A, B, C$.

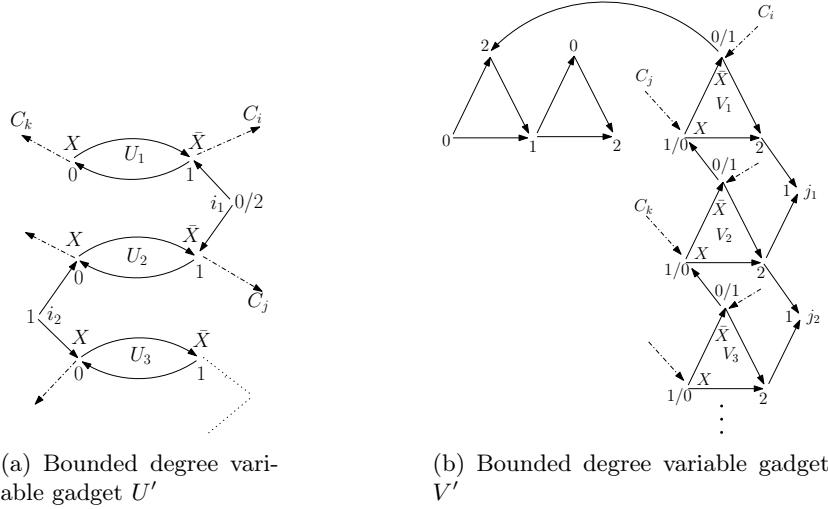


Figure 6: A - and B -coloring for $\Delta^+ \leq 2, \Delta^- \leq 2 \in \text{NP-complete}$

We now focus on input digraphs with $\Delta^+ \leq 2, \Delta^- \leq 2$. In order to prove the NP-completeness of this bounded degree setting, we need to make modifications to the gadgets used in the unbounded degree case, as the variable gadgets used in Section 2 could potentially have very high outdegree or indegree at the literal vertices. So, to keep the degree constraints satisfied, we need to construct multiple copies of each literal to join the clause gadgets of various clauses in which the literal occurs. At the same time, we have to ensure that the color assigned to the literal vertex remains same in all copies, so that we have a consistent truth value assignment.

Lemma 2. *The A -coloring problem is NP-complete, even when restricted to digraphs with $\Delta^+ \leq 2, \Delta^- \leq 2$.*

Consider the variable gadget U used in Figure 2(a). If a literal occurs x times in ϕ , then the NP-completeness construction gives the corresponding vertex in U an out-degree of $x+1$. Therefore, we shall consider a modified variable gadget U' , see Figure 6(a), which has multiple copies of each literal vertex. If X occurs x times in ϕ , and \bar{X} occurs x' times in ϕ , then we construct U' using digons U_1, U_2, \dots, U_k , where $k = \max(x, x')$.

As before, in any A -coloring each digon must be colored by 0 and 1. Because of the vertex i_1 , the first two copies of \bar{X} must be colored by the same color, 0 or 1. (Indeed, there is no vertex in A that has an edge to both 0 and 1.) Similarly, the vertex i_2 ensures that the second and third copy of X are colored by the same color. Repeating this argument we conclude that all vertices corresponding to X have the same color, and similarly for \bar{X} . This allows us to apply the previous proof to a new graph G_ϕ in which each literal vertex in U' is adjacent to at most one clause gadget W (the same clause gadget as before). Thus G_ϕ has all in- and out-degrees bounded by 2. It follows from the above remarks that each A -coloring of G_ϕ yields a satisfying truth assignment of ϕ , and it is also easy to see that any satisfying truth assignment of ϕ can be extended to an A -coloring of G_ϕ .

The analysis is again similar for B -coloring.

Lemma 3. *The B -coloring problem is NP-complete, even when restricted to digraphs with $\Delta^+ \leq 2, \Delta^- \leq 2$.*

In this case, we consider construct G_ϕ from the modified variable gadget V' in Figure 6(b), and the same clause gadget \hat{W} as before. If X occurs x times in ϕ , and \bar{X} occurs x' times in ϕ , then we

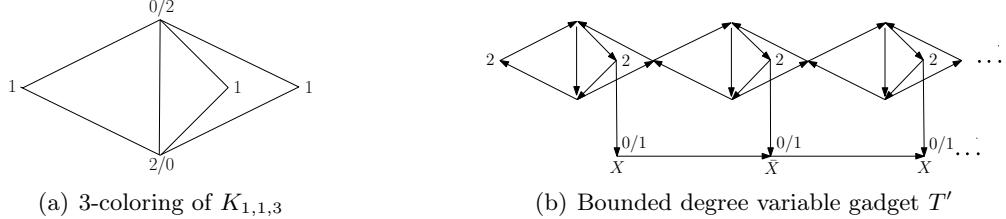


Figure 7: C -coloring for $\Delta^+ \leq 2, \Delta^- \leq 2 \in \text{NP-complete}$

construct V' using triangles V_1, V_2, \dots, V_k , where $k = \max(x, x')$, allowing for a sufficient number of copies of X and \overline{X} to keep the in- and out-degrees bounded by 2. It remains to verify that the repeated copies of X and \overline{X} must obtain the same color in any B -coloring of G_ϕ . As noted earlier, the third vertex of V_1 must be colored 2, and hence the vertex j_1 must be colored 1. This means that the third vertex of V_2 is colored 0 or 2; but 0 is not possible as the vertex has in-degree two and 20 is not an edge of B . Thus the third vertex of each V_i is 2. It now follows that in any B -coloring of V' all copies of X (and all copies of \overline{X}) receive the same color, 0 or 1. Therefore, we can repeat the previous proof of NP-completeness.

Finally, we discuss C -coloring.

Lemma 4. *The C -coloring problem is NP-complete, even when restricted to digraphs with $\Delta^+ \leq 2, \Delta^- \leq 2$.*

Here, we use the variable gadget depicted in Figure 7(b). It is an orientation of a chain of copies of $K_{1,1,3}$, and $K_{1,1,3}$ is a graph that takes the same color for every vertex of degree two. Thus assuming the leftmost vertex is colored 2 as depicted, all the vertices corresponding to X and \overline{X} are colored by 0 and 1 or 1 and 0 respectively. The overall construction is depicted in Figure 8. There is a basis triangle from which the construction emanates, and which is assumed to be colored by 0, 1, 2 as depicted (by renaming the colors if necessary). There is a *clause chain* (of copies of $K_{1,1,3}$) providing us with a sufficient number of vertices colored 1 for each clause C_i of ϕ . We use the same clause gadget W' as in the unbounded setting. Finally, there is a *variable chain* (of copies of $K_{1,1,3}$) yielding a sufficient number of vertices colored 2 for each variable X_j . The literal vertices of the clause gadgets are identified with the corresponding literals in the variable chains as before. It is easy to check that the maximum in- and out-degree is two, and that ϕ is satisfiable if and only if G_ϕ is 3-colorable.

Combining all the results for the restricted setting of bounded degree input graphs, we have our main theorem.

Theorem 4. *Deciding A -, B -, and C -colorability is in \mathbb{P} for input digraphs with $\Delta^+ \leq 2, \Delta^- \leq 1$ or $\Delta^+ \leq 1, \Delta^- \leq 2$, and is NP-complete for digraphs with $\Delta^+ \leq 2, \Delta^- \leq 2$.*

4 Conclusions

There are three smallest digraphs H with NP-complete H -coloring problems, namely A, B , and C from Figure 1. We have shown that for all three, the H -coloring problem remains NP-complete if the in- and out-degrees are bounded from above by 2, but become polynomial-time solvable if either of these degree bounds is lowered to 1. This suggests that the conjecture of Feder, Hell, and Huang, claiming that NP-complete coloring problems remain NP-complete even for bounded degree graphs provided the degree bound is high enough, may also hold for digraphs. This is the first result of this

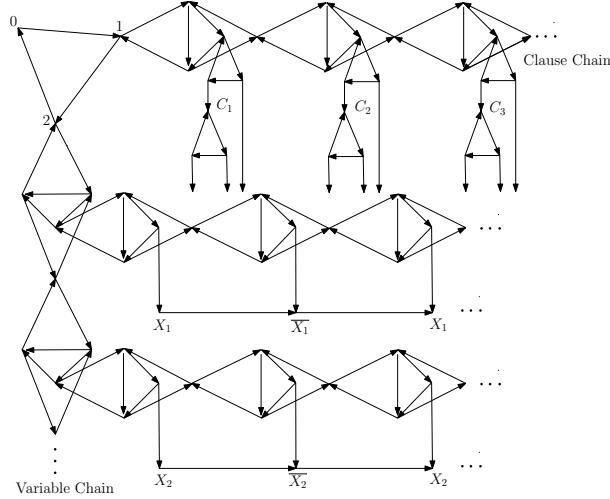


Figure 8: The overall NP-completeness construction of G_ϕ

type. Moreover, we were successful at delineating very precisely the boundary between the degree bounds that imply polynomiality and those that do not.

We conclude by noting that our proof implies the following classical result. The proof is exactly the same as for Lemma 4, except the edges of the gadgets are not oriented.

Theorem 5. [11] *The 3-coloring problem is NP-complete even when restricted to graphs with maximum degree 4.*

Acknowledgements

The authors thank the MITACS Globalink Internship Program for undergraduate students for funding the internship of the second author which made this collaboration possible. The authors would also like to acknowledge the hospitality of Simon Fraser University, and especially of the IRMACS center, where most of this was research was done, during the research internship of the second author.

References

- [1] Jørgen Bang-Jensen, Pavol Hell, and Gary MacGillivray. The complexity of colouring by semicomplete digraphs. *SIAM J. Discret. Math.*, 1(3):281–298, August 1988.
- [2] Libor Barto, Marcin Kozik, and Todd Niven. Graphs, polymorphisms and the complexity of homomorphism problems. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC ’08, pages 789–796, New York, NY, USA, 2008. ACM.

- [3] R. L. Brooks. On coloring the nodes of a network. *Mathematical Proceedings of the Cambridge Philosophical Society*, pages 194–197, 1941.
- [4] Tomás Feder, Pavol Hell, and Jing Huang. List homomorphisms of graphs with bounded degrees. *Discrete Mathematics*, 307(3-5):386–392, 2007.
- [5] Tomás Feder and Moshe Y. Vardi. The computational structure of monotone monadic snp and constraint satisfaction: A study through datalog and group theory. *SIAM J. Comput.*, 28(1):57–104, February 1999.
- [6] Michael R. Garey and David S. Johnson. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., New York, NY, USA, 1990.
- [7] P. Hell and X. Zhu. Duality and polynomial testing of tree homomorphisms. *Trans. Amer. Math. Soc*, 348:1281–1297, 1996.
- [8] Pavol Hell and Jing Huang. Counting list homomorphisms and graphs with bounded degrees. In *Discrete Math*, 2001.
- [9] Pavol Hell and Jaroslav Nešetřil. On the complexity of H-coloring. *J. Comb. Theory Ser. B*, 48(1):92–110, February 1990.
- [10] Pavol Hell and Jaroslav Nešetřil. *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [11] Ian Holyer. The NP-completeness of edge-coloring. *SIAM J. Comput.*, 10(4):718–720, 1981.
- [12] H.A. Maurer, J.H. Sudborough, and E. Welzl. On the complexity of the general coloring problem. *Information and Control*, 51(2):128 – 145, 1981.